Kinsey Reporter: Citizen Science for Sex Research

Clayton A Davis*
Center for Complex Networks and Systems Research, Indiana University School of Informatics and Computing
Bloomington, IN, USA

Julia Heiman
The Kinsey Institute for Research in Sex, Gender, and Reproduction
Bloomington, IN, USA

Erick Janssen
Institute for Family and Sexuality Studies, KU Leuven
Leuven, Belgium

Stephanie Sanders
The Kinsey Institute for Research in Sex, Gender, and Reproduction
Bloomington, IN, USA

Justin Garcia
The Kinsey Institute for Research in Sex, Gender, and Reproduction
Bloomington, IN, USA

Filippo Menczer
Center for Complex Networks and Systems Research, Indiana University School of Informatics and Computing
Bloomington, IN, USA

arXiv:1602.04878v1 [cs.CY] 16 Feb 2016

ABSTRACT

Kinsey Reporter is a global mobile app to share, explore, and visualize anonymous data about sex. Reports are submitted via smartphone, then visualized on a website or downloaded for offline analysis. In this paper we present the major features of the Kinsey Reporter citizen science platform designed to preserve the anonymity of its contributors, and preliminary data analyses that suggest questions for future research.

INTRODUCTION

The most popular sex applications available on the app marketplaces, e.g., Tinder, Grindr, and OKCupid, are great for finding sexual partners, but what happens during those and other sexual encounters? The original Kinsey Reports [5, 6] famously used face-to-face interviews to ascertain what really happens between the sheets. Adapting Kinsey's survey methodology to use web technology has provided sex researchers with additional methods for data gathering, but the increasing ubiquity of mobile computing devices presents new opportunities. Compared to traditional PCs, smartphone platforms provide location services and greater access to users in the developing world.

Kinsey Reporter (`kinseyreporter.org`) is a global mobile survey platform to share, explore, and visualize anonymous data about sex. Reports are submitted via smartphone, then visualized on the website or downloaded for offline analysis. "Citizen sex scientists" submit reports, each consisting of one or more surveys, after participating in or observing sexual activity. Surveys cover topics such as flirting, sexual activity, unwanted experience, consumption of pornography, and hormonal birth control side effects. One of the goals of the projects is to reveal how norms and behaviors surrounding these topics vary depending on geography, so location data is crucial. For example, some research indicates that hormonal birth control can differently affect women living in disparate regions of the world [13]. However the sensitive nature of these topics necessitates that reports be kept anonymous to protect both the participants and the researchers. The current implementation of Kinsey Reporter balances both of these demands by collecting no personally-identifying details and by anonymizing submission time and location with user-selectable resolution: city, province, or country.

Kinsey Reporter is a joint project of The Kinsey Institute for Research in Sex, Gender, and Reproduction and the Center for Complex Networks and Systems Research, both at Indiana University, Bloomington. Version 2 of Kinsey Reporter is currently available for iOS in the Apple App Store and for Android on Google Play. To date, about 8,300 reports have been tabulated from 32 different countries. The results can be visualized at `kinseyreporter.org/explore`.

In this paper we present the major features of the Kinsey Reporter platform, and preliminary data that suggest interesting research questions for future study.

MOTIVATION AND DESIGN

This project aims to contribute to several different bodies of research surrounding collection and analysis of sensitive and confidential data.

In order to delineate the various axes of feature space that it exists in, we have created Kinsey Reporter with an challenging set of constraints: an open-data, citizen-science platform that collects data on sensitive and private topics whilst protecting the anonymity of its users. To be clear, this represents somewhat of an extreme case as visualized in Figure 1 where

*Contact author. Email: claydavi@indiana.edu

Submitted for consideration for the PhD qualifying exam

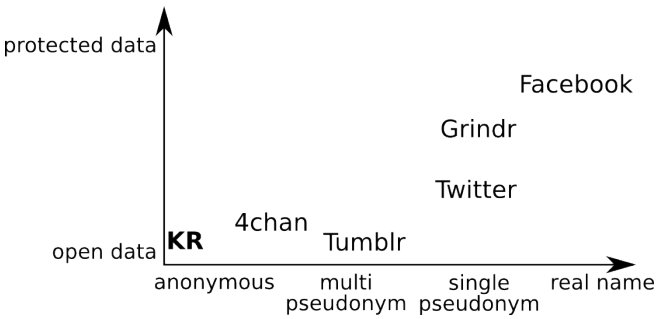

Figure 1. Two axes of feature space with some social sites located on the axes for illustration. Placement is approximate, based on typical use of the sites. KR is Kinsey Reporter.

KR is all the way in the corner. Other applications dealing with data of varied sensitivity will naturally occupy design space in between.

With this project, we hope to stimulate discussion on the pros, cons, and best practices of using citizen science platforms to collect sensitive data about private behaviors. One such question of best practices comes from noting the distinction between personal data that is "donated", as in citizen science platforms, as opposed to data that is exchanged for a service, as in the case of Foursquare or other such location-enabled apps. Are the app owners' obligations to the users different in these cases? Kinsey Reporter intentionally constrains its data collection in order to provide anonymity and protect both user and app provider from liability, but in doing so, does it also provide an example of the kind of transparency and protection from unintended use that should be expected of donated data?

Successfully interviewing subjects about sensitive topics – as did our namesake Alfred Kinsey – requires first achieving a certain level of interpersonal trust. Trust is not a binary variable however; varying degrees of trust lead to people sharing different amounts and types of information, depending on context [2]. In this light, Kinsey Reporter can be seen as a case study of how the overall platform design can contribute to the user's perception of trust.

Research in medical contexts suggests that when people are unwilling to disclose sensitive, possibly stigmatizing or embarrassing information, their fear is along two axes: fear of judgement and fear that the information will leak beyond the intended bounds [11]. We hope that Kinsey Reporter's strict anonymity protections are understood by users to protect against both of these fears. In addition, our design acknowledges the social nature of trust and privacy [2, 3]. By providing users with feedback about others' responses in the form of maps (Fig. 2) and aggregated survey results (Fig. 3), we aim to employ subtle social encouragement to participate.

ANONYMITY FEATURES

Survey topics can be highly sensitive. As such, neither the researcher nor any third-party with data access should be able to tie a report to a particular user. This implies that we can not collect demographic or identifying information about particular users.

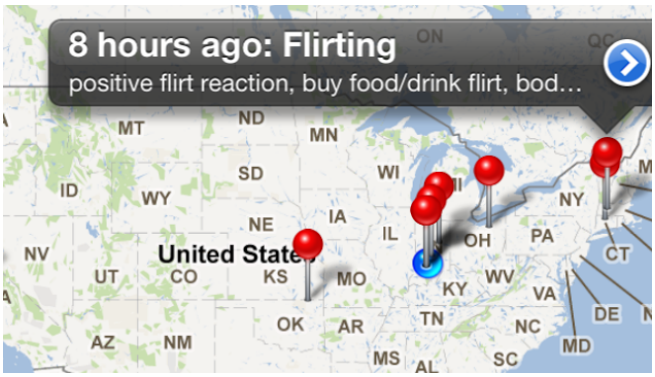

Figure 2. A cropped screenshot from the app, showing others' reports on a map.

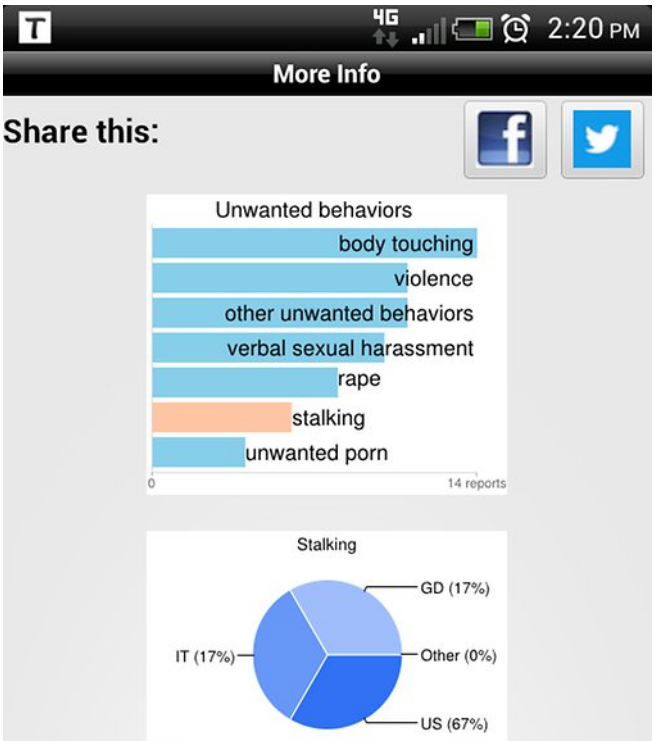

Figure 3. A screenshot from the app. After submission of a report, a user is able to explore how his or her submissions stack up to the aggregate. This user submitted the tag "stalking", thus it is highlighted in the aggregated result.

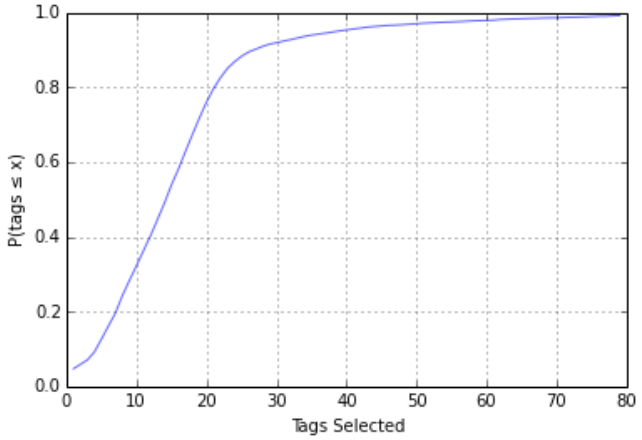

Figure 4. Empirical cumulative distribution function (CDF) of the number of tags selected per report. The mean is approximately 16 tags per response. Only 1% of reports have more than 80 tags.

A stated goal of Kinsey Reporter as a citizen science project is to provide the data open to all. As a result, pseudonymity (via usernames) becomes unacceptable: an attacker managing to de-anonymize one report should not be able to obtain the user's report history. Thus we must abandon the concept of a user identity entirely, and design for total user anonymity. These two axes of design space are illustrated in Figure 1, along with some other social web sites.

Due to these hard requirements, we cannot analyze reports grouped by user. Empirically, users selected 136 thousand answers in 8300 reports submitted, yielding about 16 selections per report. The distribution of the number of tags per response can be seen in Fig. 4. This makes it possible to examine co-occurrence by report, especially among answers to questions belonging to the same survey. One can visualize these co-occurrence relationships with the *Charts* tool at `kinseyreporter.org/explore`.

Geo-temporal indistinguishability

Suppose an attacker could observe a user interacting with Kinsey Reporter, for example in a coffee shop or other public place. If reports were posted immediately upon submission, the attacker could tie a report to the user by refreshing the report data until one matches the user's time and place of submission. To prevent this attack, reports are not made public immediately. Instead, recently-received reports sit in the database without timestamps and remain inaccessible to the public. These reports are then timestamped and made public only once a minimum number k of reports are received from a particular geographic designation. This protocol is akin to a user-adjustable, geo-temporal version of Sweeny's k-anonymity [12]. The k parameter is selected to tradeoff between data availability and user anonymity.

For a fixed k, another tradeoff exists between geographic and time resolution, as low-resolution geographic designations (e.g., country) will receive more reports in a given time than those with higher resolution (e.g., city). As a result, high geographic resolution corresponds to low temporal resolution, and vice-versa. This design is driven by an *uncertainty principle*: it should not be possible to know both the location and the time of a report with high precision, as such information might reveal too much about an individual user. On the downside, this protocol can make it difficult to compare data between rural and urban areas, due to the comparatively lower number of reports coming from a given rural municipality. It can also take a long time to accumulate k reports from some of these more remote areas, causing reports to stay in "limbo" for long periods of time.

No free text

Free-text responses lend themselves to abuse. For instance, users could submit nonsense answers to skew results, as in Mountain Dew's infamous "Dub the Dew" campaign where the number one response was "Hitler did nothing wrong" [10]. More importantly, free-text responses might be used to reveal (on purpose or not) someone's identity. This deanonymization can be on purpose, by answering a question with text like "John Doe committed sexual assault last night," or inadvertently, in the case of stylometric analysis. For example, textual analysis may be used to infer demographic information about the submitter such as age and gender, tie together multiple reports coming from the same submitter, and even possibly connect submissions to extant social media accounts [4, 7, 8].

For these reasons, users are not allowed to enter free-text responses to questions; all questions are multiple-choice. As a result, surveys must be very carefully crafted so as to avoid bias.

Unfortunately, there is little preventing users from "stuffing the ballot box" with spurious multiple-choice selections, since user anonymity precludes us from banning a user. This, however, cannot be done *en-masse*, as described next.

Secure communication

The app employs encryption and authentication in its communication with the Kinsey Reporter servers. First, public-key encryption (TLS/HTTPS) is a natural requirement to prevent eavesdropping on the network. Second, shared-key authentication is used to make it harder for an attacker script to submit a large volume of junk data to the Kinsey Reporter servers.

There are no trusted third parties. A geocoding service is used as a third-party to anonymize locations, but never receives sensitive data. The user's device only submits its location to the service, receiving the current city, state/province, and country. This lower-resolution location is then sent along with the report to Kinsey Reporter's servers.

In this case, the use of a third party for anonymizing location improves security. By ensuring that the exact location and the survey results are transmitted to separate remote servers, a hypothetical adversary needs to eavesdrop on two different encrypted connections in order to obtain unanonymized data.

SURVEY DATA

The user is presented with multiple surveys she can choose to answer, for example "Sexual Activity" or "Flirting." Associated with each of these surveys are several questions, e.g.,

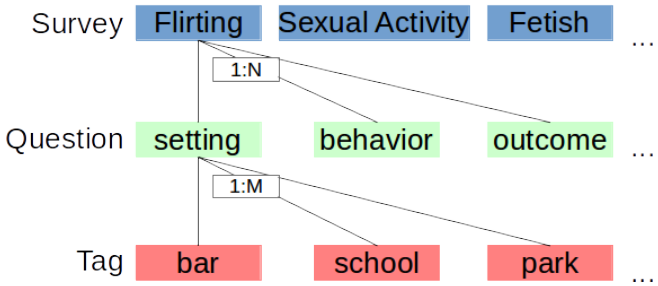

Figure 5. Data model: surveys have associated questions, which in turn contain tags.

Table 1. Countries and US states with the most responses tabulated.

Country	Count	US State	Count
USA	7138	Indiana	2907
Italy	289	California	432
Canada	237	Texas	308
Netherlands	172	Oregon	249
Great Britain	65	Ohio	224
China	63	Michigan	199
Spain	56	Arizona	195
Turkey	51	Illinois	190
Denmark	48	Kentucky	159
Australia	25	New York	147

"How many people involved?" and "Gender flirting." Likewise each question has a number of possible responses, called tags, such as "multiple partners," or "male flirting." Each of these relationships is one-to-many. The hierarchical data model is illustrated in Fig. 5.

Geography

While reports have thus far been tabulated from 32 different countries, the vast majority come from the US. The top ten countries by number of responses can be seen in Table 1 along with the top ten US states. Among US states, Indiana is over-represented due to the disproportionate amount of media coverage Kinsey Reporter receives in its home state.

While Kinsey Reporter is explicitly global in its mission, it is limited by the current lack of localization — the app and the surveys are currently only available in English. The effects of this limitation are observable in that four of the top five countries are largely English-speaking. Adding other languages will increase the potential user base of the platform and enable better international comparison of responses.

Survey items

As with the geography, the distribution of responses across surveys illustrated in Table 2 is skewed, with the Sexual Activity survey receiving more responses than all other surveys combined. In some ways this is expected simply from demographics: with sexual activity broadly defined, this survey applies to the most people. On the other hand we suspect a large amount of self-selection happening in the survey-choosing process with the more "fun" surveys receiving more responses.

Table 2. Surveys and number of tabulated responses for each. Note that the Porn and Valentine's Day surveys were released much later than the others.

Survey	Count
Sexual Activity	6605
Flirting	1161
Public Display of Affection	858
Sexual Fetish	827
Porn	528
Female Hormonal Birth Control Use and Effects	519
Unwanted Experience	306
Valentine's Day	196

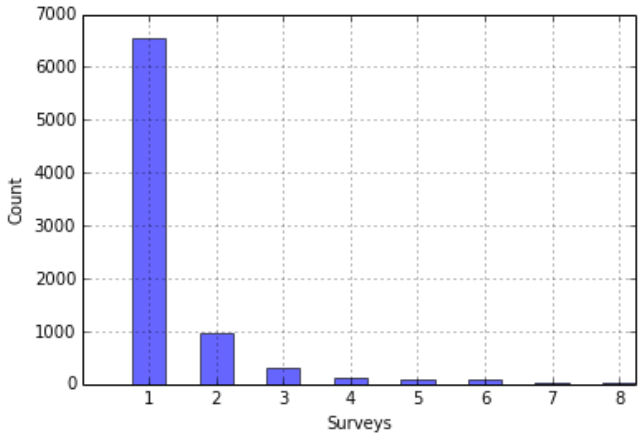

Figure 6. Histogram of number of surveys responded to per report. The majority of reports contain answers to only a single survey.

With the inability to correlate user reports across responses, co-occurrence relationships must be analyzed from each response individually. Since users choose on average 16 tags per report, we have co-occurrence data that works well when correlating responses to a single survey, e.g., sexual activity given relationship status, as both questions are from the Sexual Activity survey. However, as shown in Fig. 6, most responses contain answers from a single survey, thus somewhat limiting the ability to deduce co-occurrence relationships across surveys, e.g., sexual activity given flirting location.

Preliminary results

While Kinsey Reporter is still a work in progress, it has tabulated a fair amount of data, albeit from a rather narrow set of surveys and locations. One co-occurrence relationship from the Sexual Activity survey is shown in Fig. 7, illustrating sexual activity given relationship description. The chart shows expected relationships, such as increased incidence of oral sex amongst those in casual relationships [9], but also unexpected features such as the increased incidence of anal sex in casual encounters (392 reports). While these results are far from conclusive, these two examples combine to demonstrate the potential values of Kinsey Reporter's data: substantiating existing claims with a new source of data, and raising new questions to be further researched.

FUTURE WORK

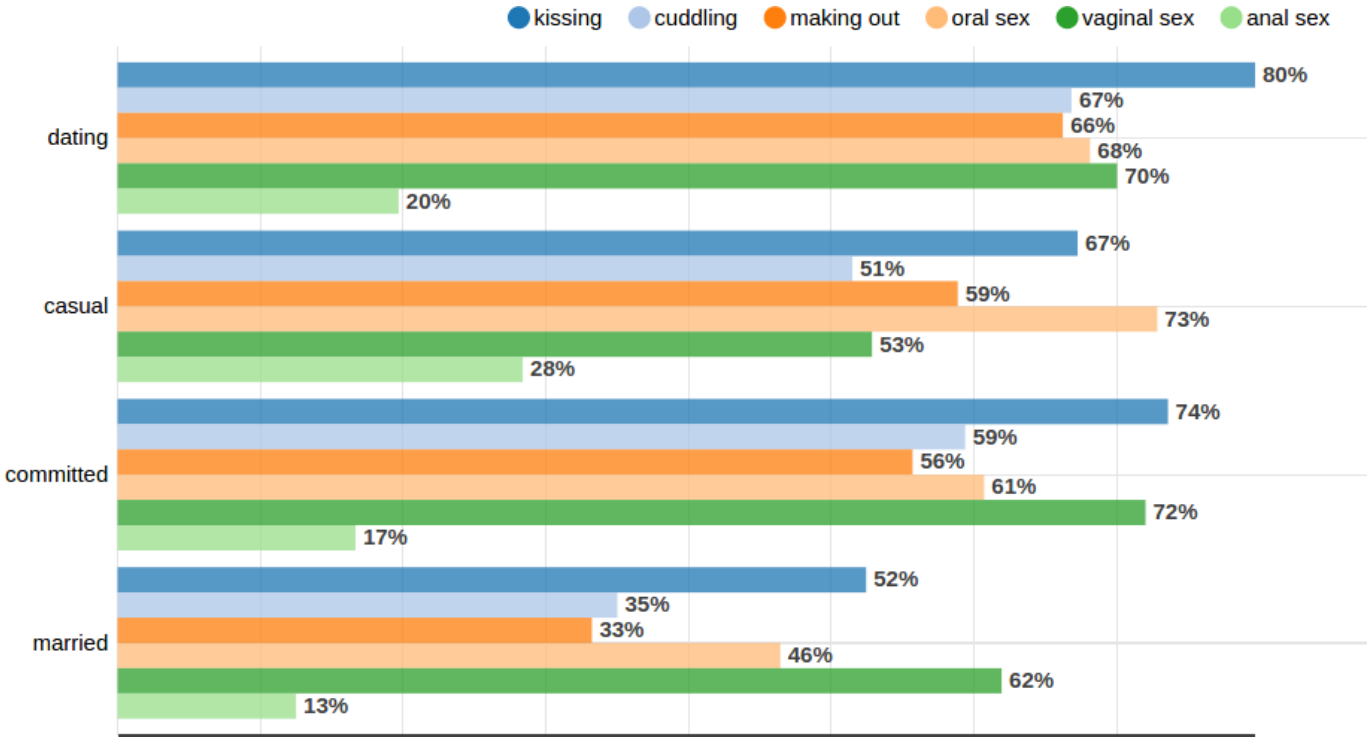

Figure 7. Sexual activity given relationship description, based on 5453 reports containing answers to both questions. Responses can include more than one activity. This chart was created with the Explore tool at `kinseyreporter.org/explore`.

Recent scholarly work has been done on privacy-preserving location-based services, including differential privacy models where the user selects the level of geographic resolution at which she is visible [1]. As mentioned previously, Kinsey Reporter's current protocol can cause reports coming from low volume areas to stay in the queue for long periods of time. We are interested in adapting some of these new techniques to preserve geo-indistinguishability while reducing the time lag before making reports public.

As mentioned previously and as seen in Fig. 6, the fact that the vast majority of submissions contain answers to only one survey somewhat limits the ability to analyze inter-survey co-occurrence relationships. A reasonable null model of user behavior might suggest half as many users submit responses to two surveys as opposed to a single submission, with the number of submissions decreasing geometrically with the number of surveys responded to, as seen in Figure 8. The overwhelming tendency to answer a single survey suggests that users may not appreciate that a report can contain responses to more than one survey (as is indeed desirable from the researchers' viewpoint).

Future collaboration with the human-computer interaction community may yield insights about how users interpret and experience the app, which in turn will allow us to analyze the data using stronger assumptions. We also welcome suggestions about necessary improvements and high-priority enhancements to attract higher volumes of data world-wide. Internationalization is also a high priority at this juncture. Our

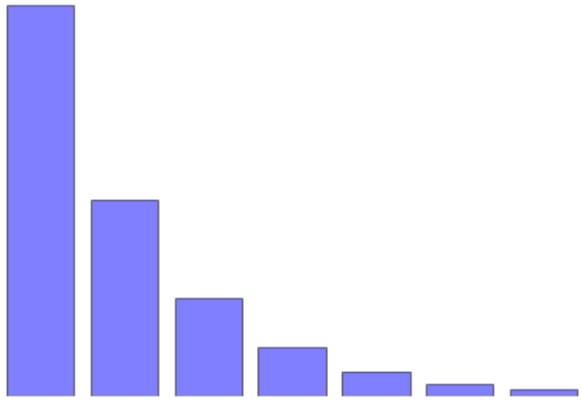

Figure 8. A null model prediction for number of surveys responded to per report. The difference with Fig. 6 may suggest a design flaw.

focus is explicitly international but our current batch of surveys are all in English.

Acknowledgements

We are grateful to Giorgio Elia, Muthu Chidambaram, Clayton Sheets, Magesh Vadivelu, and Nate Johnson for contributions to the development of the Kinsey Reporter app, and to Virginia J. Vitzthum for consultation on the surveys. Other contributors to the project are listed at `kinseyreporter.org/faq`.

REFERENCES

1. Miguel E. Andrés, Nicolás E. Bordenabe, Konstantinos Chatzikokolakis, and Catuscia Palamidessi. Geo-indistinguishability: differential privacy for location-based systems. In *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security - CCS '13*, pages 901–914, New York, New York, USA, November 2013. ACM Press.
2. Denise Anthony, David Kotz, and Tristan Henderson. Privacy in Location-Aware Computing Environments. *IEEE Pervasive Computing*, 6(4):64–72, 2007.
3. Michael Benisch, Patrick Gage Kelley, Norman Sadeh, and Lorrie Faith Cranor. Capturing location-privacy preferences: Quantifying accuracy and user-burden tradeoffs. *Personal and Ubiquitous Computing*, 15(7):679–694, 2011.
4. Fredrik Johansson, Lisa Kaati, and Amendra Shrestha. Detecting multiple aliases in social media. *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining - ASONAM '13*, pages 1004–1011, 2013.
5. Alfred Charles Kinsey, Wardell Baxter Pomeroy, and Clyde Eugene Martin. *Sexual Behavior in the Human Male*. Indiana University Press, Bloomington, 1948.
6. Alfred Charles Kinsey, Wardell Baxter Pomeroy, Clyde Eugene Martin, and Paul H. Gebhard. *Sexual Behavior in the Human Female*. Indiana University Press, Bloomington, 1953.
7. Mohammed Korayem and David J. Crandall. De-anonymizing Users Across Heterogeneous Social Computing Platforms. *Proceedings of the Seventh International AAAI Conference on Weblogs and Social Media*, 2013.
8. Claudia Peersman, Walter Daelemans, and Leona Van Vaerenbergh. Predicting age and gender in online social networks. *Proceedings of the 3rd international workshop on Search and mining usergenerated contents SMUC 11*, pages 37–44, 2011.
9. Chris Reiber and Justin R. Garcia. Hooking up: Gender differences, evolution, and pluralistic ignorance. *Evolutionary Psychology*, 8(3):390–404, 2010.
10. Everett Rosenfeld. Mountain Dew's "Dub the Dew" Online Poll Goes Horribly Wrong. *Time Newsfeed*, August 2012.
11. Timothy Stablein, Joseph Lorenzo Hall, Helen Nissenbaum, and Denise Anthony. Gay Males and Electronic Health Records: Privacy Perceptions, Age, and Negotiating Stigma. In *Proc. of the Eastern Sociological Society Annual Meeting*, 2012.
12. Latanya Sweeny. k-Anonymity: a Model for Protecting Privacy. *International Journal of Uncertainty, Puzziness and Knowledge-Based Systems*, 10(5):557–570, 2002.
13. Virginia J Vitzthum and Karin Ringheim. Hormonal Contraception and Physiology : A Research-based Theory of Discontinuation Due to Side Effects. *Studies in Family Planning*, 36(1):13–32, 2014.